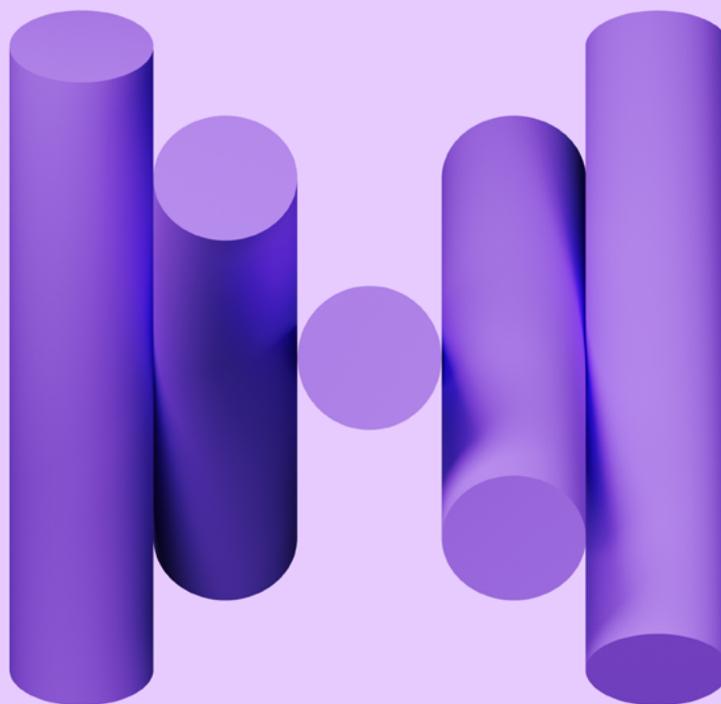


NetApp Cloud Tiering (Powered by FabricPool technology)



The intelligent way to manage your data growth

Our data explosion era

According to IDC, unstructured data will meet or exceed 175 zettabytes by 2025, a growth rate of 530% from 2018¹, with upwards of 80% of all data being unstructured². What does this mean?

Data is growing fast, as is the list of compliance requirements for long-term data retention. Along with that, systems are running out of space much faster than originally planned. On average, 80% (or more) of this data is infrequently accessed or aged (i.e., cold) data, but still resides on expensive, high-performance storage in case it's needed³.

With this overwhelming growth of data, IT admins struggle to manage their growing data stores, balancing ‘how much data I need to keep?’ with ‘how much storage can I afford?’ and ‘where should I store my data?’ Most companies can’t just delete it due to mandatory retention policies. This drives the need to find a way to store data cost-effectively, while ensuring that it’s available when needed without engineering a change in user experience, workflow, or applications.

Why does it matter

On average, most data is not accessed or read within 90 days of creation; typically, 80% or more has not been accessed in over a year. This data that is not accessed in months (or years) can tie up precious capital resources and is better suited for lower-cost object-based storage managed in the public cloud.

Hot and cold data

There are many things to consider when evaluating and determining how to manage often versus less-often used data. Assessments are normally focused on key characteristics:

- Age the data was last accessed
- Time the data was last updated
- Whether the data is considered mission or business critical
- If the data is tied to an actively used application or workload

There are two overall categories that most IT teams use to characterize their data:

- Hot data
- Cold data

Hot data is frequently accessed and actively updated, or newly created. Mission and business critical data is typically classified as hot, as is data that is tied to active applications or workloads. This data requires high-performance, low latency storage, immediately accessible upon request by a user or application. In most cases, less than 20% of all corporate data is truly hot data.

Conversely, cold data is rarely accessed, often falling into the aged data category. It is not mission or business critical and may not even be tied to an actively utilized application or workload. This data probably needs to be retained but should be considered for lower-cost object-based storage that might carry increased latency, which is understood and accepted. Customers and analysts agree that

80%, and in some cases 90%, of data falls into this category.

Storage tiering

The industry’s approach to dealing with cold data is storage tiering. A well-designed tiering approach will use policies to keep current, active, and priority data on high-performance storage, while infrequently-access, lower priority data is on lower performance, less expensive storage resources, preferably simplifying, not complicating, operations, and streamlining infrastructure management.



“Connecting the on-prem device to object storage on an AWS allows us to not only save money but also scale our on-prem environments.”

Marcus Masching
Head of IT Computing Services, Festo

[Read the cast study](#)

The answer: NetApp Cloud Tiering

NetApp’s Cloud Tiering is designed specifically to help you optimize storage costs, simplify IT operations, and streamline infrastructure.

Cloud Tiering enables you to move cold data from high-performance storage to lower-cost, object-based storage, delivering more storage for less cost. It can reduce your existing storage footprint by upwards of 80%, depending on your tiering strategy. It enables unified tiering management of all ONTAP tiering instances directly from your Cloud Manager global control plane, including any existing FabricPool licenses you may have.

Cloud Tiering continually analyzes and reports on your savings resulting from the tiering process, both real-time and historically, and project your potential savings of non-tiered volumes, enhancing your business planning and cost optimization. With the click of a button, you can initiate the Cloud Tiering instinctive wizard to tier more volumes or configure your current policies for additional savings.

Cloud Tiering simplifies your IT operations. It eliminates reengineering of application and workflows by retaining tiered data in the same namespace. Pre-defined tiering policies allows your IT resources pick the select the tiering level and quickly implement it. However, if a customized approach is deemed better, the intuitive wizards make it easy to define, apply, automate, and execute your tiering policies, with point-and-click simplicity.

Since the tiered data stays within the same namespace, in your environment, applications and workloads access the data without any reengineering, re-configuring, or re-architecting. Users have direct access to the data, eliminating the need for a data request procedure to be implemented. There is no change to application access or workflow. In other words, there is ZERO impact on applications, workloads, or users.

Cloud Tiering automatically discovers all cold data in your ONTAP on-premises and CVO clusters, down to the volume level. Cloud Tiering supports cloud-based object storage, NetApp StorageGRID, and 3rd party S3 compatible. Active IQ AI-based analytics to deliver ML-based recommended tiering actions on volumes that should be tiered or where additional economic savings are available, seamlessly through Cloud Manager.

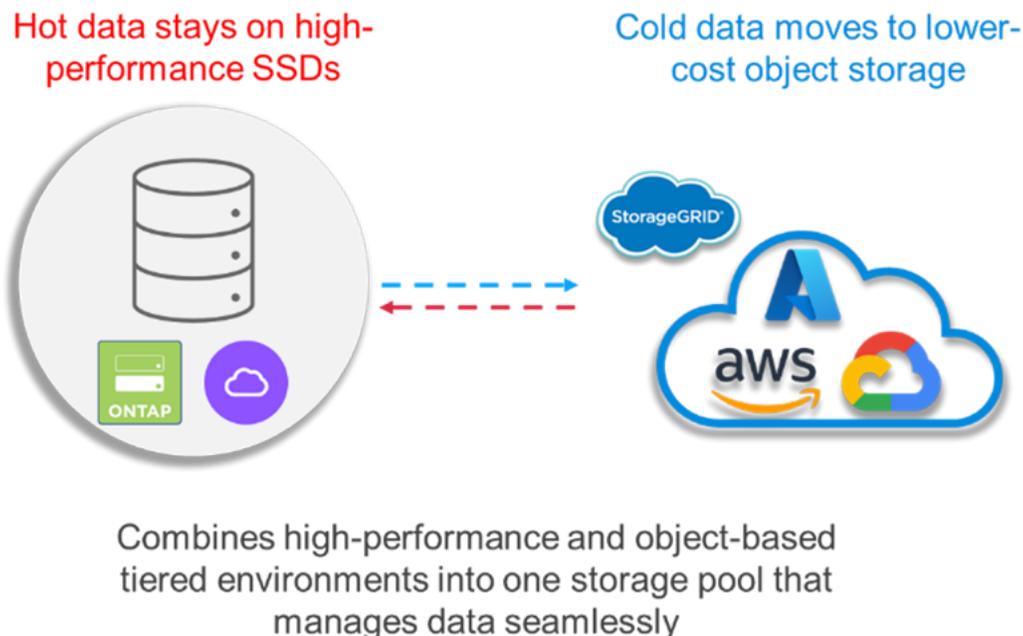
Your team can set it and forget it.

Cloud Tiering streamlines your storage infrastructure management, to extend your available capacity. Your high-performance storage capacity can be extended by 50-80%, depending on your tiering strategy. Your primary storage footprint, enabling easier management. Cloud Tiering enables high-performance, expensive storage to support more workloads without increasing capacity. Your team can implement structure storage tiers to meet your business strategy. Cloud Tiering is also an easy and safe way to start defining your cloud migration strategy.

What about FabricPool licenses?

If you have existing FabricPool licenses, you are in luck. Cloud Tiering now manages FabricPool licenses. You can manage your existing tiering instances, regardless of what systems your FabricPool licenses are licensed to, even if they came with your StorageGRID environment. Cloud Tiering even discovers your existing tiering instances to StorageGRID, delivering full visibility and management from a single pane-of-glass.

Cloud Tiering even allows you to 'float' existing FabricPool licenses to Cloud Tiering entitlements through Cloud Manager's Digital Wallet, delivering single point-of-control and increased flexibility.



How Cloud Tiering works

Cloud Tiering is part of NetApp's Cloud Volumes platform of services and enabled through NetApp Cloud Manager. The intuitive setup ensures you are running in minutes, as it automatically discovers your ONTAP systems and volumes. You can setup custom policies for tiering, or utilize the predefined options, as you specify your data cooling period for the selected volumes. Cloud Tiering starts intelligently tracking and identifying the frequently and infrequently access data blocks.

Frequently accessed data blocks are retained on primary storage, and infrequently accessed data blocks are moved to your defined object-based storage bucket. Whenever tiered data is requested, Cloud Tiering seamlessly brings it from the object tier back to the performance tier, re-heating the data.

The net-net

Cloud Tiering enables comprehensive unstructured-data management by placing your data at your right tier, in your desired location, at the right time, in accordance with your defined policies. All from a single point-of-control with unified tiering management.

You set it, forget it, and save! We take care of the rest.

Cloud Tiering's intelligent data growth management, makes tiering easy, automatic, and stress-free, delivering immediate savings, just as any tiering process should.

And best of all, it's delivered by NetApp, the cloud storage specialists.

For more information, please visit us at NetApp.com/Cloud-Services/Cloud-Tiering/ and Cloud.NetApp.com/Cloud-Tiering.

[Start your free Cloud Tiering trial.](#)



About NetApp

In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services and applications to the right people—anytime, anywhere. www.netapp.com

1 "Why unstructured data is the future of data management." July 22, 2021. VentureBeat;

"6 Predictions About Data In 2020 And The Coming Decade," January 6, 2020. Gil Press, Forbes.

2 "Five key points about unstructured data storage on-prem and cloud", Feb 04, 2021. ComputerWeekly.com.

3 "Lack of visibility is hindering hybrid data management". August 26, 2021, BlocksandFiles.com



+1 877 263 8277